arXiv:1408.1754v1 [cs.PL] 8 Aug 2014

A Partial-Order Approach to Array Content Analysis

Graeme Gange[1], Jorge A. Navas[2], Peter Schachte[1], Harald Søndergaard[1], and Peter J. Stuckey[1]

[1] Department of Computing and Information Systems,
The University of Melbourne, Victoria 3010, Australia
{gkgange,schachte,harald,pstuckey}@unimelb.edu.au
[2] NASA Ames Research Center, Moffett Field, CA 94035, USA
jorge.a.navaslaserna@nasa.gov

Abstract. We present a parametric abstract domain for array content analysis. The method maintains invariants for contiguous regions of the array, similar to the methods of Gopan, Reps and Sagiv, and of Halbwachs and Péron. However, it introduces a novel concept of an array content graph, avoiding the need for an up-front factorial partitioning step. The resulting analysis can be used with arbitrary numeric relational abstract domains; we evaluate the domain on a range of array manipulating program fragments.

1 Introduction

Most imperative programming languages offer mutable arrays. However, owing to the indirect relation between storage and retrieval, arrays are not particularly amenable to static analysis. While analysis of array *bounds* (e.g, [15,16,2]) is well studied, only recently has there been real progress in analyzing array content. Early approaches involved *array smashing* [1], where an entire array was treated as a single symbolic variable. This used the concept of *weak* updates: transfer functions for array assignment that can only weaken the previous abstract state. Weak updates generally lead to a rapid loss of precision when different segments of an array have different properties.

A significant improvement was Gopan, Reps and Sagiv's use of *array partitioning* [7] to split arrays into symbolic intervals, or segments. Partitioning is facilitated by an initial analysis of the array index expressions to determine the relative order of the indices they denote. A key idea is to distinguish segments that represent single array cells from those that represent multiple cells. This permits *strong* updates on the singleton segments. The array partitioning method selects a small set of partition variables, maintaining disjunctive information about properties which hold over all feasible total orderings of the partition variables.

Halbwachs and Péron [9] extended the approach to support relational content domains and a limited form of quantified invariants. The resulting method is precise, but has its own drawbacks. First, it requires an initial segmentation phase, where the set of partition variables is identified; and as this phase is purely syntactic, it is possible for variables to be omitted which are critical to the invariant. Second, there are exponentially many possible total orderings of the partition variables; if many partition variables are identified, the analysis may become prohibitively expensive. For example, on the $\texttt{init\_rand}_m$ family of programs shown in Figure 1, the number of partitions at the loop head follows the progression $[6, 30, 222, 2190, 27006]$ as m increases from 1 to 5 (this is discussed in more detail in Appendix C). Finally, the analysis does not support arbitrary manipulation of index variables; indices may only be incremented and decremented.

```
i_1 := 0 ; ... ; i_m := 0
x := ⋆
while(i_1 < n ∧ ... ∧ i_m < n)
    p := ⋆
    if(p < 0)
        A[i_1] := x + 1
        i_1 := i_1 + 1
    else if(p < 1)
        A[i_2] := x + 2
        i_2 := i_2 + 1
    ...
    else
        A[i_m] := x + m
        i_m := i_m + 1
```

Fig. 1: The $\texttt{init\_rand}_m$ family of program fragments.

Cousot *et al.* [3] instead maintain a single partitioning of the array, selecting a consistent totally ordered subset from a scalar variable analysis. This does not require a separate segmentation phase, saving considerable overhead; however, as it considers only a single consistent ordering (and supports only value domains) the invariants it derives are quite weak. Consider $\texttt{init\_rand}_2$, with x fixed to 0—so each element in $[0, n)$ will be assigned either 1 or 2. In this case, the relationship between i_1 and i_2 is not known, so we must select either $0 \leqslant i_1 < n$ or $0 \leqslant i_2 < n$. In either case, the desired invariant at the loop exit is lost.

An alternative approach to expressing array properties is to lift an abstract domain to quantified invariants [8]. This technique is quite general but there are two major limitations. First, it requires from the user the specification of *templates* to describe when quantifiers are introduced. Second, it is expensive, owing to the computation of *under-approximations*. For example, to join the formulas $\forall U(G \Rightarrow e)$ and $\forall U(G' \Rightarrow e')$, we must compute an under-approximation of $G \sqcup G'$, since G and G' are in negative positions, and this is prohibitively expensive for many domains.

Dillig *et al.* [6] replace strong and weak updates with *fluid* updates. Their method is a points-to and value analysis, so not relational in our sense. It builds a points-to graph where nodes represent abstract locations that include arrays qualified by index variables. Edges represent constraints on index variables that identify which concrete elements in the source location point to which concrete location in the target. A fluid update removes the dichotomy between strong and weak updates by computing first a constraint φ representing the elements that

are modified by the update. Then it adds a new points-to edge with φ (strong update) while adding the negation of φ to existing edges from the source (weak update). As φ is an over-approximation, its negation is an under-approximation and thus it would be unsound to add it directly to other edges. Instead the analysis produces *bracketing constraints* which are pairs of over- and under-approximations so that negation can be done in a sound manner. This analysis is very expressive, avoiding the large number of explicit partitions fixed a priori in [7,9]. However, the method can still be very expensive since whenever an array is accessed, the points-to edges must be modified by adding possibly disjunctive formulas.

We propose a new approach to array content analysis. We extend any existing scalar domain by introducing a pseudo-variable to refer to segments of each array, selecting index expressions as nodes in a graph, and annotating the graph edges with the properties that hold in the segments of the arrays between those index expressions. These *array content graphs* offer greater flexibility than other approaches, as they allow us to reason about properties that hold over contiguous array segments without committing to a single total ordering on index expressions, while still taking advantage of available partial ordering information. The result is an array content analysis which is fully automatic, can be used with arbitrary domains, and does not incur the up-front factorial cost of previous methods [7,9]. In particular, it can be used for relational analyses, and accounts for the possibility of array elements being related to array indices.

We base our presentation on a small control flow graph language.

Instructions	I	$\rightarrow$	$v_1 = \mathit{constant} \mid v_1 = \circ\, v_2 \mid v_1 = v_2 \diamond v_3 \mid \mathsf{A}$
Array assignments	A	$\rightarrow$	$v_1 = \mathit{arr}[v_2] \mid \mathit{arr}[v_1] = v_2$
Jumps	J	$\rightarrow$	$\mathit{if}\ (v_1 \bowtie v_2)\ \mathit{label}_1\ \mathit{label}_2 \mid \mathit{br}\ \mathit{label} \mid \mathit{error} \mid \mathit{end}$
Blocks	B	$\rightarrow$	$\mathit{label} : \mathsf{I}^*\ \mathsf{J}$
Programs	P	$\rightarrow$	B^+

Each basic block is a (possibly empty) sequence of instructions, ending in a (possibly conditional) jump. Arithmetic unary and binary operators are denoted by $\circ$ and $\diamond$ respectively, and comparison operators by $\bowtie$. We assume that there is a fixed set of arrays $\{A_1, \ldots, A_k\}$, which have global scope (and do not overlap in memory). The semantics is conventional and not discussed here. Figure 2 shows an example program in diagrammatic form.

Our analysis assumes an abstract domain $\mathcal{L} = \langle L, \sqsubseteq, \bot, \top, \sqcup, \sqcap \rangle$ for analysis over the array-free fragment of the language (obtained by leaving out A) (scalar analysis). We use this parametric domain to construct the array content analysis.

The remainder of this paper is structured as follows. Section 2 introduces the method and its underlying ideas. Section 3 discusses computational details and efficiency. The method has been evaluated experimentally; Section 4 gives a report and Section 5 concludes, suggesting further work.

2 A Graph-Based Array Content Domain

We let V and $\mathcal{A}$ be sets of scalar and array variables, respectively. A state in the concrete domain is a pair $\langle \sigma, \rho \rangle$, where $\sigma : V \mapsto \mathbb{Z}$ maps scalar variables to integer values, and $\rho : \mathcal{A} \to \mathbb{Z} \to \mathbb{Z}$ maps array cells to values.[3]

Let $V_\mathcal{A}$ denote the set of variables, and $C_\mathcal{A}$ the set of constants, which may act as segment bounds. We use an extended set of variables $V' = V_\mathcal{A} \cup C_\mathcal{A} \cup \{v^+ \mid v \in V_\mathcal{A} \cup C_\mathcal{A}\}$. The pseudo-variable v^+ represents the value $v + 1$. This allows us to talk about properties that apply to single array cells. The constant vertices are often needed; because array properties often hold over ranges bounded by a constant on one side, and array processing code often initializes a fixed set of indices before processing the rest of the array, it is not sufficient to consider only those variables and constants used directly as indices. Consider the copy program shown in Figure 2. In this case neither 0 nor N is ever used *directly* as an index, but both define boundary conditions of i, and are needed for the invariant of interest. In practice, however, $V_\mathcal{A}$ is often considerably smaller than V.

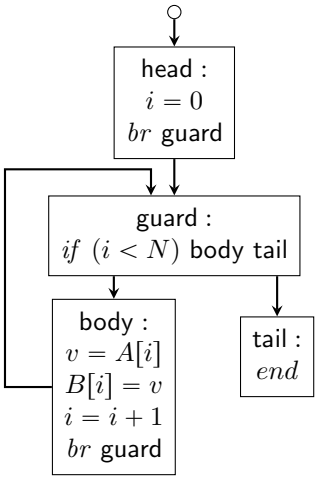

Fig. 2: Copying array A to B

We wish to relate the value of elements in an array segment to variables in the scalar domain. A state in the abstract domain has the form $\langle \varphi, \Psi \rangle$, where φ expresses scalar properties and Ψ expresses array content properties. For each array $A \in \mathcal{A}$, we allocate a corresponding variable a, a *segment* variable, which occurs only in Ψ, never in φ. Relations between scalar and segment variables are captured in Ψ. We use $U = \{a, b, \ldots\}$ to denote the set of segment variables. Sometimes, we may wish to relate the values in an array segment to the corresponding index, for example, to prove that $A[i] = x + i$ across a segment. To support this, we introduce a variable idx to represent the index of a given read. We use $U_I = \{idx\} \cup U$ to denote the augmented set of segment variables.

The analyses of the scalar domain and array contents are based on the same lattice L. We represent the program state as a pair of the scalar properties φ and the $|V'| \times |V'|$ matrix Ψ of array properties such that ψ_{ij} denotes the properties which hold for all indices in the interval $[i, j)$. In a slight abuse of notation, we use $\psi(\ell)$ to denote the formula ψ with each symbolic array variable a replaced with the corresponding array element $A[\ell]$. That is:

$$\psi(\ell) \equiv \psi\{a \mapsto A[\ell] \mid A \in \mathcal{A}\}$$

[3] For simplicity we assume arrays elements are integers. The extension to arbitrary types (that may include integers) is not difficult, as the complexity of array elements acting as indices is present in what we consider.

$\varphi = 0 \leqslant i < N$ $\quad \Psi =$

	0	i	i^+	N
0	$\bot$	$a = b$	$\top$	$\top$
i	$\bot$	$\bot$	$a = v$	$\top$
i^+	$\bot$	$\bot$	$\bot$	$\top$
N	$\bot$	$\bot$	$\bot$	$\bot$

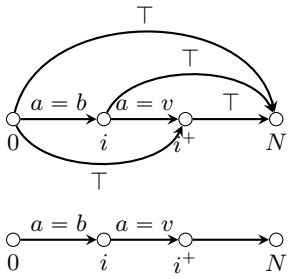

Fig. 3: Array content graph after assignment $v = A[i]$ in Figure 2. Vertices and matrix entries corresponding to 0^+ and N^+ are omitted.

Then the reading of an edge (i, ψ_{ij}, j) is: $\forall\, i \leqslant \ell < j\ .\ \psi_{ij}(\ell)$. Given some numeric abstract domain $\mathcal{L}$, a set of arrays A and scalar variables V, the array content domain $\mathcal{C}_\mathcal{L}(A, V')$ is a pair $\langle \varphi, \Psi \rangle$, where φ is a value in $\mathcal{L}$, and Ψ is a $|V'| \times |V'|$ matrix of $\mathcal{L}$-values. Assume we have a function $\mathsf{eval}_\mathcal{A}$ which constructs a new state which treats the ℓ^{th} element of each array as a scalar variable:

$$\mathsf{eval}_\mathcal{A}(\sigma, \rho, \ell) = \sigma \cup \{a \mapsto \rho(A)(\ell) \mid A \in \mathcal{A}\}$$

The concretization function γ is then defined on the components in terms of the concretization function $\gamma_\mathcal{L}$ of the scalar domain $\mathcal{L}$:

$$\gamma(\langle \varphi, \Psi \rangle) = \{\langle \sigma, \rho \rangle \mid \sigma \in \gamma_\mathcal{L}(\varphi)\} \cap \gamma_A(\Psi)$$

$$\gamma_A(\Psi) = \bigcap_{i,j \in V'} \{\langle \sigma, \rho \rangle \mid \forall_{\sigma(i) \leqslant \ell < \sigma(j)} \mathsf{eval}_\mathcal{A}(\sigma, \rho, \ell) \in \gamma_\mathcal{L}(\psi_{ij})\}$$

Notice that there are no constraints on ρ in the first equation; ρ can be any array variable assignment of type $\mathcal{A} \to \mathbb{Z} \to \mathbb{Z}$. As any value in the content domain $\mathcal{C}_\mathcal{L}(A, V')$ is the Cartesian product of a fixed set of elements of $\mathcal{L}$, $\mathcal{C}_\mathcal{L}(A, V')$ also forms a lattice, and possesses all the corresponding fixed point properties.

For each edge ψ_{ij}, we can assume the corresponding interval is non-empty; that is, $[\![i < j]\!] \in \psi_{ij}$. Note that the edge from i to i^+ has no such constraint, since $i < i{+}1$ is always true. Since the interval $[i{+}1, i)$ is clearly empty, $\psi_{i^+ i} = \bot$.

Given the state representation, we can take the join of two abstract states by a piecewise application of the join from $\mathcal{L}$:

$$\langle \varphi^1, \Psi^1 \rangle \sqcup \langle \varphi^2, \Psi^2 \rangle = \left\langle \varphi^1 \sqcup \varphi^2, \begin{bmatrix} \psi^1_{11} \sqcup \psi^2_{11} & \dots & \psi^1_{1n} \sqcup \psi^2_{1n} \\ \vdots & \ddots & \vdots \\ \psi^1_{n1} \sqcup \psi^2_{n1} & \dots & \psi^1_{nn} \sqcup \psi^2_{nn} \end{bmatrix} \right\rangle$$

We can compute the meet $\sqcap$ analogously. To see how the analysis works, consider again the program in Figure 2. Figure 3 shows the abstract state immediately after executing $v = A[i]$. The array content information is given by the matrix

Ψ of array properties. The *array content graph* shown upper right is really just a way of visualizing the matrix. Note that infeasible edges, those labelled $\bot$, are omitted. In fact, we shall usually show only the "transitive reduction" of the array content graph, so that an edge ik whose value is given by $\psi_{ik} = \bigsqcap_j(\psi_{ij} \sqcup \psi_{jk})$ is omitted. We depict edges representing $\top$ without a label. This leads to the shorthand graph in Figure 3's lower right.

2.1 Normalization of Abstract States

Given a set C of constraints of the form $e \in [i, j) \Rightarrow \psi_{ij}$, we wish to normalize the state by computing the strongest consequences of C, still in that form.

The critical observation is this: $\forall x, y \;.\; x < y \models \forall z.(x < z) \vee (z < y)$. That is, any property that holds over both $[x, z)$ and $[z, y)$ must also hold over the range $[x, y)$. To compute the strongest consequences, then, we must compute the greatest fixed point of a rewrite system derived from the set of inequalities: $\forall i, j, k \;.\; \psi_{ij} \sqsubseteq \psi_{ik} \sqcup \psi_{kj}$.

It is tempting to try to compute the fixed point using the obvious rewrite system:

$$\forall i, j, k \;.\; \psi_{ij} = \psi_{ij} \sqcap (\psi_{ik} \sqcup \psi_{kj})$$

But, as we shall see, this formulation does not guarantee termination for all useful domains.

Mohri [12] describes the *algebraic shortest path* problem, where the operations $(\sqcap, \sqcup)$ form a semiring in the domain of edge weights. This is very close to what we need, as every distributive lattice forms a bounded semiring. But "numeric" domains used in static analysis generally fail to be distributive, so we cannot use Mohri's framework directly.

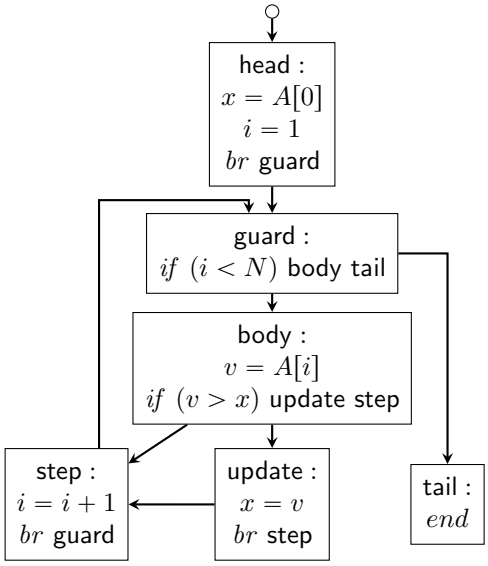

Fig. 4: Find maximum value in array A

Example 1. Consider the array-maximum program from Figure 4. Figure 5(1) shows the program state just after $v = A[i]$ is executed (scalar constraints on the left and array constraints on the right). On the branch with $v \leqslant x$, we simply add the constraint to the scalar domain, resulting in (2a). If $v > x$, we add the constraint to the scalar domain (2b), then update x with v, resulting in (2b'). Observe that, in both cases, we can only discover the relationship between a and x indirectly via v. If we do not push scalar relations into the edge properties, the underlined invariants are lost. The final result is shown at (3).

If the statement $x = v$ in update was instead replaced by $x = A[i]$, we would first have to lift the invariant $x = v$ from the singleton property ψ_{ii^+} to the

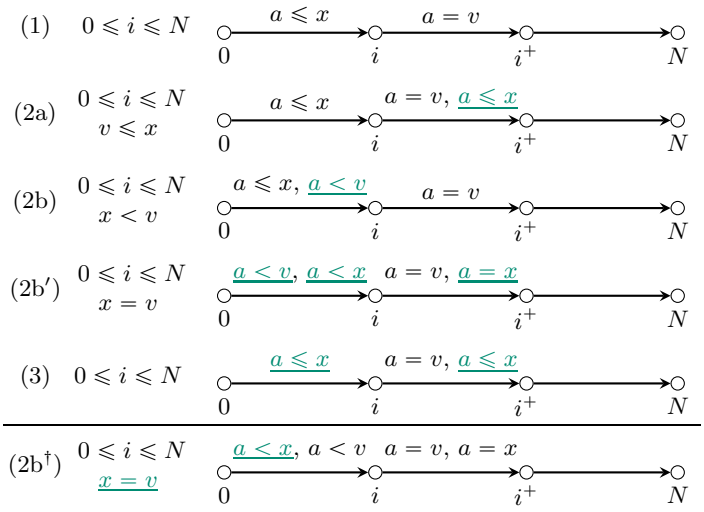
(1) $0 \leqslant i \leqslant N$
$a \leqslant x$
$a = v$
(2a) $0 \leqslant i \leqslant N$ $v \leqslant x$
$a \leqslant x$
$a = v$, $\underline{a \leqslant x}$
(2b) $0 \leqslant i \leqslant N$ $x < v$
$a \leqslant x$, $\underline{a < v}$
$a = v$
(2b′) $0 \leqslant i \leqslant N$ $x = v$
$\underline{a < v}$, $\underline{a < x}$
$a = v$, $\underline{a = x}$
(3) $0 \leqslant i \leqslant N$
$\underline{a \leqslant x}$
$a = v$, $\underline{a \leqslant x}$
(2b†) $0 \leqslant i \leqslant N$ $\underline{x = v}$
$\underline{a < x}$, $a < v$
$a = v$, $a = x$
0
i
i^+
N

Fig. 5: Array maximum example

scalar domain. We then push this property into the segment ψ_{0i}, allowing us to derive $a < x$. The state in Figure 5(2b†) shows the program state if $x = v$ were to be replaced with $x = A[i]$. ∎

This illustrates that it is not sufficient to simply compute the transitive closure of Ψ; we must also lift properties from Ψ out to the scalar domain. For a fully reduced state $\langle \varphi, \Psi \rangle$, the following properties must be satisfied for all i, j, k:

1. The graph of segments must be internally consistent: $\psi_{ij} \sqsubseteq \psi_{ik} \sqcup \psi_{kj}$
2. Segment properties are consistent with the scalar domain: $\psi_{ij} \sqsubseteq \varphi \sqcap [\![i < j]\!]$
3. For each non-empty segment, the scalar domain must be consistent with the scalar properties of that segment: $\varphi \models i < j \Rightarrow \varphi \sqsubseteq \exists U_I \,.\, \psi_{ij}$

Notice that we only propagate constraints to the scalar component from segments that are known to be non-empty. If we tried to propagate information from all segments, we would incorrectly derive $\bot$ as soon as any segment was determined to be empty. One solution is to simply apply the three rules until a fixed point is reached. This is guaranteed to compute the fully reduced state. However, while this direct construction is conceptually clean, it suffers from some pragmatic issues relating to both termination and efficiency, as we shall see.

2.2 Termination

The normalization process is not guaranteed to terminate for arbitrary lattices.

Example 2. Assume the analysis uses convex polyhedra. Consider the state in Figure 6(a). Any fixed point will satisfy the properties $A \sqsubseteq B \sqcup C$ and $B \sqsubseteq A \sqcup C$.

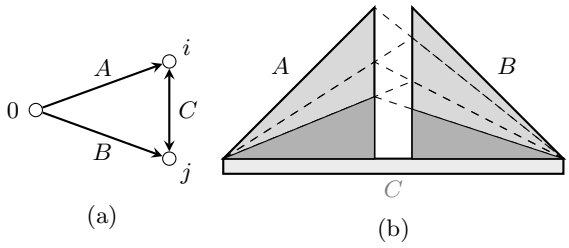

Fig. 6: At a fixed point, we have $A \sqsubseteq B \sqcup C$ and $B \sqsubseteq A \sqcup C$. The regions A and B will be progressively reduced, indefinitely.

Let A, B and C be the gray regions shown in Figure 6(b)—the intention is that A shares a line segment with C, as does B. Assume we start by exploiting $B \sqsubseteq A \sqcup C$. We compute $A \sqcup C$, yielding the polygon given by the topmost dashed line. This allows us to trim the top portion of B. We then compute $B \sqcup C$, and trim the top-left region of A. However, now A has changed, so we re-compute $A \sqcup C$ and again reduce B. This process asymptotically approaches the greatest fixed point $A = A \sqcap C,\ B = B \sqcap C$. ∎

If we modify the equations in Example 2 slightly, to $A \sqsubseteq (A \sqcap B) \sqcup (A \sqcap C)$ (still a valid approximation of the concrete state), convergence is immediate.

The fixed point process will clearly terminate for the interval domain, as the possible interval end-points are drawn from the initial set. We can also show that it is guaranteed to terminate for both octagons [11] and convex polyhedra [4]; proofs are given in Appendix A. Unfortunately, we do not yet have a more general characterisation of the lattices for which termination is (or is not) guaranteed.

2.3 Abstract Transfer Functions

In this section, we describe the abstract transfer functions necessary to perform array content analysis on the language described in Section 1.

Variable assignment. The effect of a scalar assignment $x = f$ on an abstract state follows the behaviour of the underlying domain L. We first project out the previous value of x (assuming x does not occur in f) then introduce the new constraint into the scalar domain. However, when we project x from our scalar domain, we must also update all incoming and outgoing edges of x (and x^+). This becomes: $\langle \varphi, \Psi \rangle [\![\mathtt{x} = \mathtt{f}]\!] = \langle \varphi [\![\mathtt{x} = \mathtt{f}]\!], \Psi' \rangle$ where Ψ' is given by, for all p, q:

$$\psi'_{pq} = \begin{cases} \top & \textbf{if } p \in \{x, x^+\} \vee q \in \{x, x^+\} \\ \exists x \,.\, \psi_{pq} & \textbf{otherwise} \end{cases}$$

The assumption that f is free of x is not always well founded; however, we can always transform the program so that it *is* the case: $\mathtt{x} = \mathtt{f} \Rightarrow \mathtt{x_t} = \mathtt{x};\ \mathtt{x} = \mathtt{f}[\mathtt{x}/\mathtt{x_t}]$.

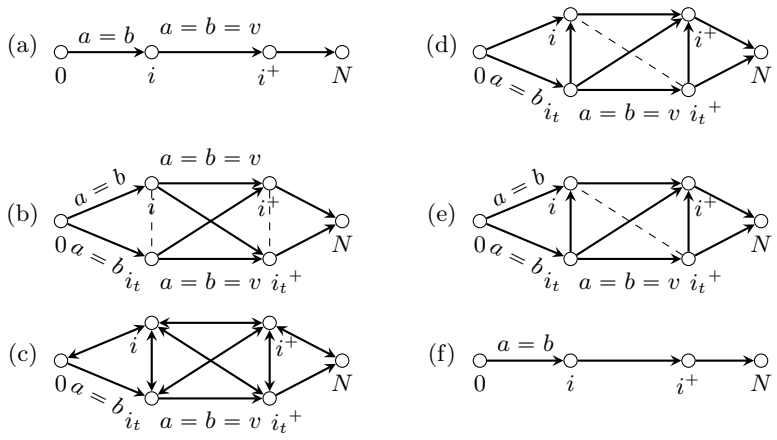

Fig. 7: Array content graph during analysis of the program given in Figure 2, while executing $i = i + 1$. The dashed edges indicate equality (in reality they represent two edges each labelled $\bot$).

Example 3. Consider the array-copy program given in Figure 2, immediately before the assignment `i = i + 1`. The state is shown in Figure 7 (a).

We must first introduce a new variable i_t to hold the prior value of i, transforming $i = i + 1$ to $i_t = i;\ i = i_t + 1$. The normalized graph after the $i_t = i$ statement is shown in Figure 7 (b). To handle the $i = i_t + 1$ statement we must eliminate the annotations on edges corresponding to i and i^+, resulting in state (c). Note that the edge $\psi_{0i_t{}^+}$, omitted from the diagram, has annotation $a = b$. State (d) completes the handling of the $i = i_t + 1$ statement, introducing the new value of i into the scalar domain. The scalar domain discovers that ${i_t}^+ = i$ (the dashed edge). When normalizing state (d), $\psi_{i_t{}^+ i} = \bot$, so the rule $\psi_{0i} \sqsubseteq \psi_{0i_t{}^+} \sqcup \psi_{i_t{}^+ i}$ results in ψ_{0i} becoming $a = b$, the desired invariant (e), which after projecting out i_t and i_t^+ gives (f). ∎

Array reads. An array read $x = A[i]$ is relatively simple. As for standard variable assignment, we must existentially quantify the variable x. But instead of introducing a relation into the scalar domain, we add the constraint $x = a$ to the singleton segment ψ_{ii^+}. This transfer function may be formulated as $\langle \varphi, \Psi \rangle [\![\mathtt{x = A[i]}]\!] = \langle \exists x.\varphi, \Psi' \rangle$ where Ψ' is given by, for all p, q:

$$\psi'_{pq} = \begin{cases} \psi_{pq} [\![\mathtt{a = x}]\!] & \textbf{if } p = i, q = i^+ \\ \exists x \,.\, \psi_{pq} & \textbf{otherwise} \end{cases}$$

Normalization handles the consequences for the scalar part.

Array writes. When we store a value into an array at index i, we update the corresponding edge property ψ_{ii^+}. However, this is not sufficient, as the singleton

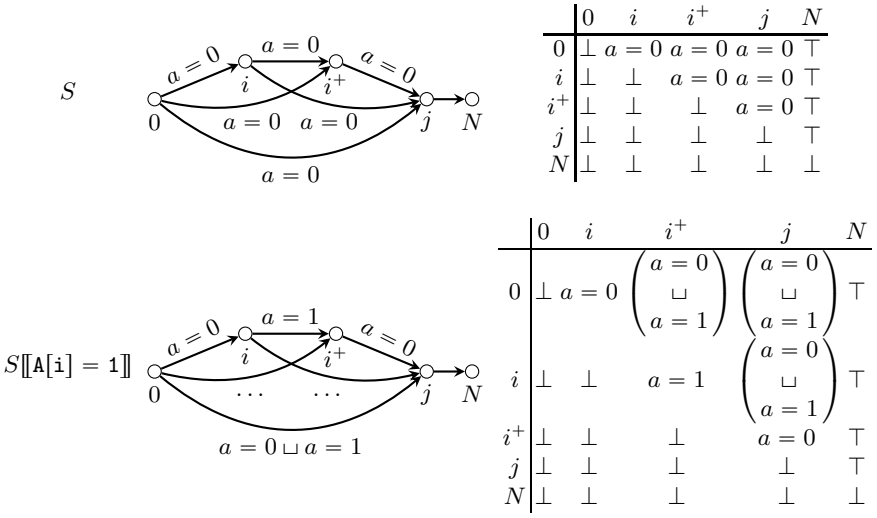

	0	i	i^+	j	N
0	$\bot$	$a=0$	$a=0$	$a=0$	$\top$
i	$\bot$	$\bot$	$a=0$	$a=0$	$\top$
i^+	$\bot$	$\bot$	$\bot$	$a=0$	$\top$
j	$\bot$	$\bot$	$\bot$	$\bot$	$\top$
N	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$

	0	i	i^+	j	N
0	$\bot$	$a=0$	$\begin{pmatrix} a=0 \\ \sqcup \\ a=1 \end{pmatrix}$	$\begin{pmatrix} a=0 \\ \sqcup \\ a=1 \end{pmatrix}$	$\top$
i	$\bot$	$\bot$	$a=1$	$\begin{pmatrix} a=0 \\ \sqcup \\ a=1 \end{pmatrix}$	$\top$
i^+	$\bot$	$\bot$	$\bot$	$a=0$	$\top$
j	$\bot$	$\bot$	$\bot$	$\bot$	$\top$
N	$\bot$	$\bot$	$\bot$	$\bot$	$\bot$

Fig. 8: In state S, we know that all elements between 0 and j have been initialized to 0. When we evaluate $\texttt{A[i]} = \texttt{1}$, we update the edge ψ_{ii^+}. However, i may also be covered by the edge ψ_{0j}. For these potentially overlapping edges, we must perform a *weak update*, taking the join of the previous value with the new.

i may be covered by other edges. As with previous analyses, we distinguish between *strong updates* where all elements in a (generally, singleton) segment are updated with a given property, and *weak updates* [1] where some elements of a segment *may* be updated. An edge ψ_{pq} must be updated if $p \leqslant i \leqslant q \wedge p < q$ is possible in the current state. This is possible if and only if the edges ψ_{pi^+} and ψ_{iq} are both feasible.

Consider the array state illustrated in Figure 8, where $0 \leqslant i \leqslant j \leqslant N$ and the first j elements have been initialized to 0. When we store 1 at $A[i]$, we update the singleton ψ_{ii^+} with the property $a = 1$. However, there are other segments that may contain $A[i]$. The segments ψ_{0i^+}, ψ_{ij} and ψ_{0j} are all consistent with index i. In this case, they must all be weakly updated (annotations on ψ_{0i^+} and ψ_{ij} have been omitted; they are identical to the annotation on ψ_{0j}).

We can formulate this as $\langle \varphi, \Psi \rangle \llbracket \texttt{A[i]} = \texttt{f} \rrbracket = \langle \varphi, \Psi' \rangle$, where Ψ' is given by, for all p, q:

$$\psi'_{pq} = \begin{cases} \psi_{pq} \llbracket \texttt{a} = \texttt{f} \rrbracket & \textbf{if } p = i, q = i^+ \\ \psi_{pq} & \textbf{if } \varphi \models (p > i \vee q < i^+) \\ \psi_{pq} \sqcup \psi_{pq} \llbracket \texttt{a} = \texttt{f} \rrbracket & \textbf{otherwise} \end{cases}$$

Notice that if we have some other variable j such that $\varphi \models i = j$, we will initially perform only a weak update of the segment ψ_{jj^+}. However, the normalization procedure will then enforce consistency between ψ_{ii^+} and ψ_{jj^+}.

3 Improving Efficiency through Relaxation

When computing the strongest matrix entries, we perform a substantial amount of redundant work. We include all constraints φ from the scalar domain in each matrix entry ψ_{ij}, and these constraints will be processed during each step in the shortest-path computation. This is not ideal, as many are irrelevant to the content properties, and abstract domain operations are often proportional to the number of constraints. Hence we want to construct some relaxation $\ddot{\psi}_{ij}$ of ψ_{ij} that discards irrelevant scalar properties. We shall use $\odot$ to denote this relaxation operation (that is, $\ddot{\psi}_{ij} = \varphi \odot \psi_{ij}$). At a minimum, it must satisfy:

$$\psi_{ij} \sqsubseteq \varphi \sqcap \ddot{\psi}_{ij} = \varphi \sqcap (\varphi \odot \psi_{ij})$$

We want to make $\varphi \sqcap \ddot{\psi}_{ij}$ as close to ψ_{ij} as possible, while keeping the representation of $\ddot{\psi}_{ij}$ concise. This is very similar to the process of constraint abduction [10]. However, even with a relaxation $\ddot{\Psi}$ such that $\varphi \sqcap \ddot{\psi}_{ij} \equiv \psi_{ij}$, we may still lose relevant information.

Example 4. Consider a program state (using the domain of `octagons`) with scalar property $\varphi = [\![x < y]\!]$, and segment properties $\psi_{ik} = \varphi \sqcap [\![i < k \wedge x < a]\!]$, $\psi_{kj} = \varphi \sqcap [\![k < j \wedge y = a]\!]$. Computing the value of ψ_{ij} gives the expected $\psi_{ij} = \varphi \sqcap [\![i < j \wedge x < a]\!]$. However, although the relaxations

$$\ddot{\psi}_{ik} = [\![i < k \wedge x < a]\!] \quad \text{and} \quad \ddot{\psi}_{kj} = [\![k < j \wedge y = a]\!]$$

are exact, computing ψ_{ij} as before yields:

$$[\![i < j \wedge i < k \wedge x < a]\!] \sqcup [\![i < j \wedge k < j \wedge y = a]\!] = [\![i < j]\!] \quad \blacksquare$$

We could avoid this loss of information by conjoining the scalar part during each step of the fixed point computation

$$\ddot{\psi}_{ij} \sqsubseteq \varphi \odot ((\varphi \sqcap \ddot{\psi}_{ij} \sqcap \ddot{\psi}_{ik}) \sqcup (\varphi \sqcap \ddot{\psi}_{ij} \sqcap \ddot{\psi}_{kj}))$$

but while this avoids the loss of information, it also defeats the original goal of reducing computation cost. Instead we define a more conservative $\odot$ operation which maintains enough additional information to retain properties of interest.

Example 5. Consider again the analysis that was performed in Example 4, but now with a modified $\odot$ operation which gives us $\ddot{\psi}_{ik} = [\![i < k \wedge x < a]\!]$ and $\ddot{\psi}_{kj} = [\![k < j \wedge y = a \wedge x < a]\!]$. In this case, when we compute $\ddot{\psi}_{ij}$, we get: $\ddot{\psi}_{ij} = [\![i < j \wedge i < k \wedge x < a]\!] \sqcup [\![i < j \wedge k < j \wedge y = a \wedge x < a]\!] = [\![i < j \wedge x < a]\!]$. This maintains the property of interest, without repeatedly conjoining with φ during the fixed point process. $\blacksquare$

Another observation from Examples 4 and 5 is that while the boundary constraints for each edge (such as $[\![i < j]\!]$ for ψ_{ij}) are not implied by φ, they are typically irrelevant to the segment properties, unless:

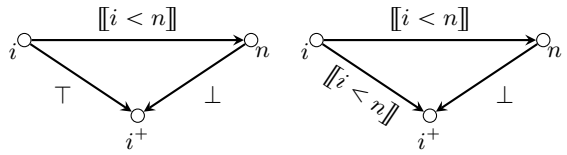

Fig. 9: Solving this set of constraints, we strengthen ψ_{ii^+} and derive $[\![i < n]\!]$. If we use $\odot$, we have $\ddot{\psi}_{in} = \top$, so we fail to strengthen $\ddot{\psi}_{ii^+}$.

- $\varphi \sqcap [\![i < j]\!] = \bot$, in which case the edge must be empty, or
- the array content $A[\ell]$ is some function of the index ℓ
- $\psi_{ji^+} = \bot$ (or $j = k^+$ for some k, and $\psi_{ki} = \bot$)

This is particularly troublesome for domains that explicitly store the transitive closure of constraints, as we then expend substantial computation maintaining consequences of $[\![i < j]\!]$, which are largely irrelevant, and lost during the join.

If we choose an operator $\odot$ which discards the consequences of edge boundaries, we must take particular care not to lose information in the third case mentioned. Consider the state shown in Figure 9, where the edge from n to i^+ is infeasible. Computing the original fixed point, we obtain $[\![i < n]\!]$ at ψ_{ii^+}, which is then lifted out to the scalar domain. However, $[\![i < n]\!]$ is obviously implied by $\varphi \sqcap [\![i < n]\!]$, so will typically be discarded by $\odot$. The property $[\![i < n]\!]$ is not obtained at ψ_{ii^+}, so is never lifted out to the scalar domain. We must, therefore, add the following case to the normalization rules given in Section 2.1:

4. $\varphi \sqsubseteq [\![i \geqslant j]\!]$ **if** $\psi_{i,j} = \bot$

This ensures that any scalar properties resulting from infeasible segments are included in the scalar domain. Further, since many elements of $\ddot{\Psi}$ may be $\top$, the relaxation operator allows us to take advantage of sparse matrix representations.

Unfortunately we are not aware of any existing, general operations suitable for computing $\ddot{\Psi}$; as they must consider both underlying lattice and the characteristics of the implementation. We complete this section by outlining suitable relaxation operators for difference-bound matrices (DBMs) [5] (and `octagons`). Relaxations for `polyhedra` can be found in the appendix, and it should not be particularly difficult to define analogous operators for alternative domains.

A value in the `DBM` (or `octagon`) domain consists of a set of constraints $[\![v_i - v_j \leqslant k]\!]$ (or $[\![\pm v_i \pm v_j \leqslant k]\!]$ for `octagon`). We can construct a relaxation of ψ_{ij} by computing the transitive closure $\psi^\star_{ij}$, and discarding any constraints implied by $\varphi \sqcap [\![i < j]\!]$:

$$\varphi \odot \psi_{ij} = \{c \mid c \in \psi^\star_{ij},\ \varphi \sqcap [\![i < j]\!] \nvdash c\}$$

If the abstract states are stored in closed form, we can simply collect the constraints of ψ_{ij} not appearing in $\varphi \sqcap [\![i < j]\!]$. Or we can avoid performing many implication tests by instead collecting the constraints involving variables in U.

4 Experimental Evaluation

We have implemented the analysis in SPARCOLYZER, a prototype array content analyser for the language described in Section 1. SPARCOLYZER is implemented in OCAML, using the FIXPOINT library[4]. For the underlying domain, we implemented the DBM domain in C++, customized for operating on sparse graphs. Experiments were performed on a 3.0GHz Core 2 Duo with 4Gb ram running Ubuntu Linux 12.04. The set $V_{\mathcal{A}}$ of segment bounds were pre-computed using a simple data-flow analysis to collect all variables which may (possibly indirectly) be involved in the computation of an array index.

We tested SPARCOLYZER on a number of array manipulation program fragments. Most of these were taken from Halbwachs and Péron [9]. We added several additional fragments that illustrate interesting properties of the analysis, including members of the `init_rand` family discussed in Section 1 (Figure 1).

Computation time for each instance is given in Table 1. For instances taken from [9], we include the original reported runtimes, although these are not directly comparable, as the experiments in [9] were performed on a slower machine (Core2 Duo 1.6 GHz, with 2MB of RAM) using the domain of *difference bound matrices with disequalities (`dDBM`)* [13].[5] The implementation from [9] has not been available and we have not tried to reconstruct it.

Table 1 compares the runtimes of two variants of our content domain to the approach of Halbwachs and Péron [9] (the hp_{08} column). The `naive` variant uses the direct implementation, where the matrix is represented as a $|V'| \times |V'|$ array, and a copy of the scalar domain is stored in each matrix entry. The `sparse` variant stores, for each row and column, a set of non-$\top$ entries so that the normalization, $\sqcup$ and $\sqcap$ operations do not need to process entries that will definitely remain $\top$, and computes the fixed point on the relaxed matrix $\ddot{\Psi}$, rather than directly on Ψ. `sparse` uses the simple relaxation step of discarding all constraints not involving some array variable $A \in U$. Note that `sparse` still iterates over all $|V'| \times |V'|$ elements when changes to the scalar domain occur, as a change to the scalar domain may affect any matrix element.

In cases where there are very few partitions—either because there are very few index variables, or they are highly constrained—we expect the partition-based methods to be faster (as they do not need to compute closure over transitive edges). The performance of `naive` is comparable to that of [9] on instances with few partitions, and it improves substantially on more complex instances. `sparse` is faster yet, sometimes by several orders of magnitude, and still finds the desired invariant in all but two cases.

It is interesting to compare the behaviour of `sentinel` and `first_nonnull`. These programs superficially appear quite similar; in both cases, we set up an 'end-of-array' marker, then scan the array to find a particular element. However, the invariants necessary to prove the desired properties are quite different. In

[4] `http://pop-art.inrialpes.fr/people/bjeannet/bjeannet-forge/fixpoint/`

[5] Performance of the domains should be roughly equivalent, as in the absence of explicit disequalities, `dDBM` behaves identically to `DBM`.

program	naive	sparse	hp$_{08}$ [9]
init	0.11	0.02	
init_offset	0.28	0.07	0.05
init_rand$_2$	2.48	0.14	
init_rand$_3$	9.59	0.64	
init_rand$_4$	31.14	1.98	
init_rand$_5$	80.58	4.96	
arraymax	0.13	<0.01	0.10
copy	0.13	<0.01	0.02
partition_hoare	1.51	0.06	
partition_hp08	3.50	0.14†	22.87
sentinel	0.14	<0.01	0.21
first_nonnull	0.60†	0.01†	2.25

Table 1: Analysis times in seconds. Instances where we were unable to prove the desired invariant are marked with †.

```
x := A[0]
i := 1
j := n − 1
while(i ⩽ j)
    if(A[i] < x)
        A[i − 1] := A[i]
        i := i + 1
    else
        while(j ⩾ i ∧ A[j] ⩾ x)
            j := j − 1                (⋆)
        if(j > i)
            A[i − 1] := A[j]
            A[j] := A[i]
            i := i + 1
            j := j − 1
A[i − 1] := x
```

(a) partition_hp08

```
x := A[0]
i := 0
j := n − 1
while(i ⩽ j)
    while(A[i] ⩽ x ∧ i ⩽ j)
        i := i + 1
    while(A[j] > x ∧ i ⩽ j)
        j := j − 1
    if(i ⩽ j)
        t := A[j]
        A[j] := A[i]
        A[i] := t
```

(b) partition_hoare

Fig. 10: Quicksort partitioning: (a) as done in [9]; and (b) a la Hoare

the case of first_nonnull, we require:

$$(s = n \wedge \forall e \in [0, n)\ .\ A[e] \neq 0) \vee (s < n \wedge A[s] = 0 \wedge \forall e \in [0, s)\ .\ A[e] \neq 0)$$

This can be expressed using the approaches of Gopan *et al.* and Halbwachs and Péron [7,9], as they store a separate invariant for each total ordering amongst the partition variables. Our approach, however, cannot handle such disjunctive reasoning, so the segment property quickly reaches $\top$.

Consider partition_hp08 (the variant of the Quicksort partition step given in [9], shown in Figure 10(a)). As the imperative source language we use does

allow loads inside conditionals, such reads are hoisted outside the corresponding loops; for example, the loop marked ($\star$) is transformed as shown below. At the point marked ($\star$), it is possible to determine that $[\![j < i]\!] \Rightarrow [\![A[j] < x]\!]$. Thus it is easy to show that $[\![A[j] < x]\!]$ holds at the loop exit. In the hoisted version, `naive` method can prove the invariant successfully, because the property $[\![j < i]\!] \Rightarrow [\![ej < x]\!]$ is derived for the edge ψ_{ji}. When we exit the loop with $[\![j < i]\!]$, this property gets extracted to the scalar domain, and we get $[\![ej < x]\!] \wedge [\![ej = a]\!]$ at $\psi_{jj'}$. When using the `sparse` method, however, the property on ψ_{ji} is discarded, as it involves only scalar variables, so the invariant is lost.

$ej := A[j]$
while$(j \geqslant i \wedge ej \geqslant x)$
  $j := j - 1$
  $ej := A[j]$

If we were to use the original version without hoisting, we would be unable to prove the invariant using either method, as we cannot express $[\![j < i]\!] \Rightarrow [\![A[j] < x]\!]$ directly.

However, our method easily proves the standard version (`partition_hoare`; Figure 10(b)) correct, whether or not the reads are hoisted.

5 Conclusion and Future Work

We have described a new approach to automatic discovery of array properties, inspired by algebraic shortest-path algorithms. This approach retains much of the expressiveness of the partitioning methods of [7] and [9], but avoids the need for syntax dependence and an up-front factorial partitioning step. The method can successfully derive invariants for a range of interesting array program fragments, and is substantially faster than partitioning-based approaches with even modest numbers of index variables.

Several improvements could be made to the performance of the analysis. The current implementation does not take advantage of liveness information, and maintains entries in the content graph for all variables in $V_{\mathcal{A}}$ at each step. Clearly, performance could be improved by eliminating non-live variables from the matrix.

Algorithms which maintain shortest path information can often be improved by storing only the *transitive reduction* of the graph. As our domains are not distributive, it is non-trivial to determine whether a given edge must occur in the transitive reduction; however, it would be worth investigating whether maintaining the transitive reduction would prove beneficial.

References

1. B. Blanchet, P. Cousot, R. Cousot, J. Feret, L. Mauborgne, A. Miné, D. Monniaux, and X. Rival. Design and implementation of a special-purpose static program analyzer for safety-critical real-time embedded software. In T. Mogensen, D. Schmidt, and I. H. Sudborough, editors, *The Essence of Computation: Complexity, Analysis, Transformation*, volume 2566 of *LNCS*, pages 85–108. Springer, 2002.

2. R. Bodík, R. Gupta, and V. Sarkar. ABCD: Eliminating array bounds checks on demand. In *ACM Symposium on Programming Language Design and Implementation (PLDI'00)*, pages 321–333. ACM Press, 2000.
3. P. Cousot, R. Cousot, and F. Logozzo. A parametric segmentation functor for fully automatic and scalable array content analysis. In *Proceedings of the 38th ACM Symposium on Principles of Programming Languages*, pages 105–118. ACM Press, 2011.
4. P. Cousot and N. Halbwachs. Automatic discovery of linear constraints among variables of a program. In *Proceedings of the Fifth ACM Symposium on Principles of Programming Languages*, pages 84–97. ACM Press, 1978.
5. D. L. Dill. Timing assumptions and verification of finite-state concurrent systems. In *Automatic Verification Methods for Finite State Systems*, volume 407 of *LNCS*, pages 179–212. Springer, 1989.
6. I. Dillig, T. Dillig, and A. Aiken. Fluid updates: Beyond strong vs. weak updates. In A. D. Gordon, editor, *Proceedings of the 19th European Symposium on Programming*, volume 6012 of *LNCS*, pages 246–266. Springer, 2010.
7. D. Gopan, T. Reps, and M. Sagiv. A framework for numeric analysis of array operations. In *Proceedings of the 32nd ACM Symposium on Principles of Programming Languages*, pages 338–350. ACM Press, 2005.
8. S. Gulwani, B. McCloskey, and A. Tiwari. Lifting abstract interpreters to quantified logical domains. In *Proceedings of the 35th ACM Symposium on Principles of Programming Languages*, pages 235–246. ACM Press, 2008.
9. N. Halbwachs and M. Péron. Discovering properties about arrays in simple programs. In *ACM Symposium on Programming Language Design and Implementation (PLDI'08)*, pages 339–348. ACM Press, 2008.
10. M. J. Maher. Herbrand constraint abduction. In *Proceedings of the 20th IEEE Symposium on Logic in Computer Science*, pages 397–406. IEEE Comp. Soc., 2005.
11. A. Miné. The octagon abstract domain. *Higher-Order and Symbolic Computation*, 19(1):31–100, 2006.
12. M. Mohri. Semiring frameworks and algorithms for shortest-distance problems. *Journal of Automata, Languages and Combinatorics*, 7(3):321–350, 2002.
13. M. Péron and N. Halbwachs. An abstract domain extending difference-bound matrices with disequality constraints. In B. Cook and A. Podelski, editors, *Verification, Model Checking and Abstract Interpretation*, volume 4349 of *LNCS*, pages 268–282. Springer, 2007.
14. A. Schrijver. *Theory of Linear and Integer Programming*. Wiley, 1986.
15. N. Suzuki and K. Ishihata. Implementation of an array bound checker. In *Proceedings of the Fourth ACM Symposium on Principles of Programming Languages*, pages 132–143. ACM Press, 1977.
16. H. Xi and F. Pfenning. Eliminating array bound checking through dependent types. In *ACM Symposium on Programming Language Design and Implementation (PLDI'98)*, pages 249–257. ACM Press, 1998.

Appendix A: Termination Proofs

Theorem 1. *The shortest-path computation terminates for the octagon domain.*

Proof. Consider an initial set of abstract states $X = [x_1, \ldots, x_n]$, and a system of inequalities of the form $x_i \sqsubseteq (x_i \sqcap x_j) \sqcup (x_i \sqcap x_k)$. Let $\mathcal{BP}(x)$ denote the bounding hyperplanes of x. We then have: $\mathcal{BP}((x \sqcap y) \sqcup (x \sqcap z)) \subseteq \mathcal{BP}(x \sqcap y) \cup \mathcal{BP}(x \sqcap z)$. Each time one of the equations is evaluated, each bounding hyperplane is an element of: $\bigcup_{X' \subseteq X} \mathcal{BP}(\bigsqcap X')$ As the initial set of bounding hyperplanes is finite, and each iteration must tighten at least one bounding plane, the tightening process must eventually terminate; in fact, the number of descending steps is bounded by: $|X| \times |\bigcup_{x \in X} \mathcal{BP}(x)|$ □

In the case of octagons, projection cannot introduce new bounding hyperplanes, so the addition of the propagation rules ($\varphi \sqsubseteq \exists\ U_I\ .\ \varphi_{ii^+}$ and $\varphi_{ij} \sqsubseteq \psi$) does not affect termination.

Theorem 2. *The shortest-path computation terminates for the convex polyhedron domain.*

Proof. We can prove termination for convex polyhedra in a similar fashion as for octagons. Consider an initial set of abstract states $X = [x_1, \ldots, x_n]$, and a system of inequalities of the form $x_i \sqsubseteq (x_i \sqcap x_j) \sqcup (x_i \sqcap x_k)$. By the polyhedron decomposition theorem (see, e.g. [14]), any polyhedron may be generated by a finite set of points P, and a set of rays R. Let R_V denote the set of rays of unit length in the direction of variable v (that is, a vector with v^{th} component 1 or -1, and all other components 0), for each $v \in V$. Given polyhedra $x = \langle P_x, R_x \rangle$ and $y = \langle P_y, R_y \rangle$, the results of the operations of interest have the following properties:

$$\begin{aligned} x \sqcap y &= \langle P', R' \rangle \ \textbf{where}\ \ R' \subseteq R_x \cup R_y \\ x \sqcup y &= \langle P', R' \rangle \ \textbf{where}\ \ P' \subseteq P_x \cup P_y,\ R' \subseteq R_x \cup R_y \\ \exists\ V\ .\ x &= \langle P', R' \rangle \ \textbf{where}\ \ P' \subseteq P_x,\ R' \subseteq R_x \cup R_V \end{aligned}$$

None of these operations can introduce rays not in $\mathcal{R} = R_V \cup \bigcup_{x \in X} R_x$. Additional extreme points can only be introduced during the application of $\sqcap$. During each fixed point iteration, exactly one of three cases must occur:

1. All abstract values in X remain the same.
2. The set of rays for some $x \in X$ changes.
3. All the sets of rays remain the same, and the set of extreme points for some $x \in X$ changes.

In case 1, we terminate. As the set of rays is restricted to $\mathcal{R}$, and each iteration is strictly descending, case 2 can only occur finitely many times. Assuming the set of rays remains fixed, every introduced extreme point must be some element of $\bigcup_{X' \subseteq X} \{P \mid \langle P, R \rangle = \bigcap_{x \in X'} x\}$. As this set is finite, and each step is descending, this can only occur finitely many times without case 2 occurring. As case 2 must always occur after a bounded number of steps, and can only occur finitely many times, the fixed point process must eventually terminate. □

We conclude that this process will terminate on the most commonly used relational numeric domains. In the case of polyhedra, it is worth noting that Theorem 2 does not provide any bounds on the *coefficients* of hyperplanes in the resulting polyhedra. In some cases, the coefficients may grow quite large before converging, which can cause problems for domains implemented with fixed-precision machine arithmetic.

Appendix B: Relaxation for `polyhedra`

The relaxation algorithm for `polyhedra` follows the same intuition as that for `octagons`; we wish to collect the transitive closure of ψ_{ij}, then discard anything implied separately by $\varphi \sqcap [\![i < j]\!]$.

However, the `polyhedra` domain provides two difficulties: computing the transitive closure of a set of linear constraints is non-trivial (as the domain is typically stored as a minimal set of generators and hyperplanes [4]); and often is a bad idea in general: as the constraints do not have bounded arity, there may be exponentially more constraints in the transitive closure than in the original problem. Instead, we collect a set of constraints that is sufficient to reconstruct the constraints that, had we computed the transitive closure, would have been kept.

Given constraints $c_1 = [\![k_1^\intercal x \geqslant m_1]\!]$ and $c_2 = [\![k_2^\intercal x \geqslant m2]\!]$, we can construct a new constraint (that is not implied separately by either c_1 or c_2) by resolution if there is some matched pair of coefficients k_{1v}, k_{2v}, such that $k_{1v} > 0$ and $k_{2v} < 0$. We then construct a new constraint $c_1 = [\![(k_1 + \frac{k_{1v}}{k_{2v}} k_2)^\intercal x \geqslant m_1 + \frac{k_{1v}}{k_{2v}} m_2]\!]$, which has v^{th} coefficient 0.

Example 6. Consider the constraints $c_1 = [\![x + y \geqslant 7]\!]$, $c_2 = [\![z - 2y \geqslant 2]\!]$, and $c_3 = [\![w + 2y \geqslant 3]\!]$. c_1 and c_2 may be resolved, as c_1 contains the term y, and c_2 contains $-2y$. This yields $c_{12} = [\![x + \frac{z}{2} \geqslant 8]\!]$. We cannot, however, construct any new constraints by combining c_1 with c_3. ∎

Rather than computing the transitive closure explicitly, given an initial set of *interesting* constraints $\ddot{C}$ and other constraints C, we find all those constraints that are *resolvable* with those in $\ddot{C}$ (taking into account the direction of previous resolution steps), and add them to $\ddot{C}$. We then continue this process until no further resolvable constraints are found:

$$\mathsf{trans}^\star(R, \ddot{C}, C) = \begin{cases} \mathsf{trans}^\star(R \cup \{(\mathsf{sign}(k_{v'}), v')\}, \ddot{C} \cup \{c\}, C) \\ \quad \textbf{if } \exists\ (s, v) \in R,\ c = [\![k^\intercal x \geqslant m]\!] \in C,\ v' \in x \\ \quad \textbf{s.t. } v \neq v',\ k_v s < 0,\ ((\mathsf{sign}(k_{v'}), v') \notin R \vee c \notin \ddot{C} \\ \ddot{C} \textbf{ otherwise} \end{cases}$$

$$\mathsf{trans}(\ddot{C}, C) = \mathsf{trans}^\star(\{\mathsf{sign}(k_v, v) \mid [\![k^\intercal x \geqslant m]\!] \in \ddot{C},\ v \in x\}, \ddot{C}, C)$$

Given a set of linear constraints C defining a polyhedron, we can construct the initial set $\ddot{C}$ with the elements of C that are not implied by $\varphi \sqcap [\![i < j]\!]$, then

compute the relaxation with $\mathsf{trans}(\ddot{C}, C \setminus \ddot{C})$. As with `octagons`, we can avoid performing implication tests by instead initializing $\ddot{C}$ with those constraints containing variables in $V_{\mathcal{A}}$.

Example 7. Consider a constraint of interest $c_0 = [\![a - y \geqslant 0]\!]$, with additional constraints $C = \{c_1 = [\![y - z \geqslant 0]\!], c_2 = [\![z + w \geqslant 0]\!], c_3 = [\![x - y \geqslant 0]\!]\}$. Initially, we have $\ddot{C} = \{[\![a - y \geqslant 0]\!]\}$, and $R = \{(+, a), (-, y)\}$. We can resolve with $[\![y - z \geqslant 0]\!]$, since $(-, y) \in R$, but $(-, z) \notin R$. During the second step, we include c_2, adding $(-, w)$ to R. At this point, we have $R = \{(+, a), (-, y), (-, z), (+, w)\}$, and $\ddot{C} = \{c_0, c_1, c_2\}$, and we cannot add anything to either R or $\ddot{C}$; so we return the current value of $\ddot{C}$. ∎

Appendix C: Partitions for `init_rand`$_m$

The large number of partitions required for the `init_rand` family is not necessarily obvious. Assuming n is non-negative, we must distinguish between the case where the array is empty (thus $n = 0$) or non-empty ($n > 0$). This gives us two base orderings: $[\{0, n\}]$, and $[\{0\}, \{n\}]$. In these descriptions, sets denote equivalence classes, and equivalence classes are listed in increasing order.

If $n = 0$, there is only one possible value for i_1: $[\{0, i, n\}, \{i^+\}]$. Otherwise, we must distinguish all the possible relations between 0, i and n. The resulting orderings are as follows:

$[\{0, i_1, n\}, \{i_1^+\}]$	$[\{0\}, \{i_1, n\}\{i_1^+\}]$
$[\{0, i_1\}, \{i_1^+, n\}]$	$[\{0\}, \{i_1\}, \{i_1^+, n\}]$
$[\{0, i_1\}, \{i_1^+\}, \{n\}]$	$[\{0\}, \{i_1\}, \{i_1^+\}, \{n\}]$

When we construct the partitions for $m = 2$, we introduce i_2 into all feasible locations in each of the partitions for $m = 1$.

In cases where we must explicitly distinguish the 0^{th} element, this progression grows substantially faster, with $m = 1, \ldots, 5$ yielding $[9, 45, 333, 3285, 40509]$.